



Release B Planning and Scheduling Hardware Selection

Gary Roth

groth@eos.hitc.com

19 April 1996

Overview



- **Design Drivers**
- **Sizing Analysis**
- **Configuration**
- **Backup and Recovery**
- **Scalability**
- **Capacity Break-points**

Design Drivers



Level 3 and 4 Requirements:

EOSD 3700	ECS functions shall have an operational availability of .96 at a minimum and a MDT of four hours or less, unless otherwise specified.
PGS 1270	Design must accommodate expansion by up to a factor of 3 without design changes, and up to a factor of 10 without major design changes.

Design Drivers (con't.)



Software Implementation Drivers:

- **AutoSys job volume:** driven by number of PGE activations per day.
- **AutoSys GUI displays:** driven by COTS software implementation.
- **PDPS database transaction load:** driven by the creation of a 30 day plan.



Hardware Sizing Analysis

Scheduling Drivers

- **AutoSys Job Volume**
 - Dependent on number of PGE activations, number of job activations per PGE, and reprocessing
- **AutoSys GUIs RAM requirements**
 - Dependent on the number of jobs for which status is displayed
- **Backup and Recovery In Support of Scheduling**
 - Second machine required to run AutoSys Shadow Event Processor which performs automatic backup and recovery
- **Growth**
 - Support increases in production processing and reprocessing

Planning Drivers

- **PDPS (Planning and Data Processing Subsystem) Database**
 - Contains critical data (PGE attributes, etc.)
 - Peak transaction load occurs when creating a 30 day plan
- **Backup and Recovery In Support of PDPS Database**
 - Second machine required to run Sybase replication server for PDPS database which performs automatic backup and recovery
 - Maintains duplicate database to ensure safety of critical PGE information
- **Growth**
 - Support increases in production processing and reprocessing



Planning/Scheduling Assumptions

Assumptions (based on 2/96 technical baseline for Epoch k, 3rd qtr '99):

- Working with MODIS to limit PGE activations to 4,000 at GSFC
- GSFC and LaRC operate 24 hours per day 7 days a week

DAAC <u>Site</u>	# of PGE <u>Activations</u>	# of jobs <u>per PGE</u>	Repro- <u>cessing</u>	Total # of <u>jobs/day</u>
GSFC	4,004	4	2	32,032
LaRC	778	4	3	9,336



AutoSys GUIs RAM Analysis

Assumptions (at Epoch k):

- Four GUIs are active on one platform (worst case)
- Status of 32,032 GSFC or 9,336 LaRC jobs being displayed
- Each GUI requires 20,000 kbytes of RAM as a base plus 5.5 kbytes for each job

$[20,000 \text{ kbytes} + (5.5 \text{ kbytes/job} \times 32,032 \text{ jobs})] \times 4 \text{ GUIs} = 785 \text{ Mbytes of RAM at GSFC}$

$[20,000 \text{ kbytes} + (5.5 \text{ kbytes/job} \times 9,336 \text{ jobs})] \times 4 \text{ GUIs} = 286 \text{ Mbytes of RAM at LaRC}$

PDPS DBMS Peak Transaction Load Analysis



Assumptions (at Epoch k, 3rd qtr '99):

- **The Planning Workbench function of creating a 30 day plan imposes the greatest peak transaction load on the PDPS database**
- **A 30 day plan requires the input of 45 days worth of data**
- **Transaction loads double due to Sybase replication (database backup)**

GSFC peak PDPS DBMS transaction load is 9,009,000 per 30 day plan creation

LaRC peak PDPS DBMS transaction load is 1,751,000 per 30 day plan creation



Hardware Assumptions

- AutoSys customer experience has indicated AutoSys can sustain 45 jobs per minute while running on a 75 MHz Sun Sparc 20/712 2-CPU machine.
- It is estimated that our throughput will be 29 jobs per minute (42,120 jobs/day) when using the AutoSys high availability option (shadow event processor).
- It is estimated a 4-CPU Sun UltraServer would have a throughput of 187,200 jobs/day.
- A 75 MHz Sun Sparc 20/712 2-CPU machine is rated at 305 transactions per second (TPS). It would take approximately 2 hours to process LaRC's 30 day plan creation peak load of 1,751,000 transactions ($1,751,000 / (305 \text{ tps} \times 60 \text{ sec/min})$).
- It is estimated that a 4-CPU Sun UltraServer TPS rating equals 1,360. It would take approximately 2 hours to process GSFC's 30 day plan creation peak load of 9,009,000 transactions.



Hardware Assumptions (con't)

	<u>Required</u>	<u>Provided Capacity</u>
GSFC		
<ul style="list-style-type: none">• Job Throughput• Time to process a 30 day plan (9,009,000 DBMS transactions)	32,032 jobs/day reasonable	187,200 jobs/day 111 min.
LaRC		
<ul style="list-style-type: none">• Job Throughput• Time to process a 30 day plan (1,751,000 DBMS transactions)	9,336 jobs/day reasonable	42,120 jobs/day 96 min.

GSFC Planning & Scheduling Hardware Configuration



Scheduling

- 4 CPUs
- 2 GB of RAM per machine (GUIs distributed between multiple X-terminals)
- 21 GB of disk (4.2 GB disk required; 200 MB for AutoSys database + 4 GB for PDPS backup database)
- Two Fast Wide SCSI controller cards per machine


Planning

- 4 CPUs
- 2 GB of RAM per machine
- 21 GB of disk (4.2 GB of disk required; 4 GB for PDPS database + 200 MB for AutoSys backup database)
- Two Fast Wide SCSI controller cards per machine

GSFC Hardware Configuration (and EDC and NSIDC)




Planning



(Production Planner)

Sun 20/71



(Resource Planner)


Sun 20/71

PDPS DBMS Server:

5 x 2.1 GB



5 x 2.1 GB



- PDPS database
- AutoSys Shadow Event Processor
- 4 CPUs
- 2 GB RAM


Scheduling

Scheduling Server:

5 x 2.1 GB



5 x 2.1 GB



- AutoSys Event Processor
- Replicated PDPS DBMS
- 4 CPUs
- 2 GB RAM

X terminals



(Production Monitors)

- AutoSys GUIs

Sun UltraServers

LaRC Planning & Scheduling Hardware Configuration




Scheduling

- 2 CPUs
- 512 MB of RAM
- 12.6 GB of disk (56 MB for AutoSys database)
- Two Fast Wide SCSI controller cards per machine

Planning

- 2 CPUs
- 512 MB of RAM
- 12.6 GB of disk (780MB for PDPS database)
- Two Fast Wide SCSI controller cards per machine




LaRC Hardware Configuration (and JPL)

Planning

Scheduling


PDPS DBMS Server:



- PDPS DBMS
- Planning Workbench
- 2 CPUs
- 512 MB RAM

Sun 20/712
(Resource Planner)


PDPS DBMS Server:



- Replicated PDPS DBMS
- Planning Workbench
- 2 CPUs
- 512 MB RAM

Sun 20/712
(Production Planner)


Scheduling Server:



- AutoSys Event Processor
- 2 CPUs
- 512 MB RAM

Sun 20/712

Scheduling Server:



- AutoSys Shadow Event Processor
- 2 CPUs
- 512 MB RAM

Sun 20/712



(Production Monitors)

X terminals

- AutoSys GUIs



Other Considerations



- **External disk used; RAID disk mirroring is not needed with AutoSys high availability option and Sybase replication of the PDPS database**
- **2.1 GB disks used instead of 4.2 GB disk in order to provide faster disk access and to provide more “spindles” which allows distribution of Sybase data, log, and index files**
- **Scheduling X-terminals: a pair of screens is required to properly display the status of production jobs to the user**
- **Design supports SSI&T access to planning workbench functions**



Backup and Recovery

GSFC

- **Scheduling Server machine or disk failure:**
 - AutoSys Shadow Event Processor on the PDPS DBMS Server automatically continues to execute the schedule
- **PDPS DBMS Server or disk failure:**
 - replicated PDPS DBMS on Scheduling Server automatically continues with PDPS DBMS transactions
- **Planning Workstation failure:**
 - use second workstation for both Production Planning and Resource Planning

LaRC

- **Scheduling Server machine or disk failure:**
 - AutoSys Shadow Event Processor on the other Scheduling Server machine automatically continues to execute the schedule
- **PDPS DBMS Server or disk failure:**
 - replicated PDPS DBMS on the other PDPS DBMS Server machine automatically continues with PDPS DBMS transactions

Scalability



GSFC

- **Sun 8-slot UltraServer can use up to 12 CPUs (from 4 CPUs)**
 - Sun 12-slot UltraServer can use up to 16 CPUs
- **Disk storage can increase to 50.4 GB per machine (from 21 GB)**
 - External SparcStorage Arrays can be added to increase disk capacity

LaRC

- **Sun 20/712 CPUs can be upgraded from 75 MHz to 150 MHz**
- **Disk storage can increase to 50.4 GB per machine (from 12.6 GB)**
 - External SparcStorage Arrays can be added to increase disk capacity



GSFC Break-point Analysis

Current hardware fully populated (12 CPU UltraServers):

Scheduling		Current	Fully
	<u>Required</u>	<u>Design</u>	<u>Populated</u>
• Job Throughput	32,032	187,200	561,600
• RAM	1.6 GB	2 GB	6 GB
• Disk	4.2 GB	21 GB	50.4 GB
Planning			
• Time to process a 30 day plan	reasonable	111 min.	37 min.
• RAM	1.6 GB	2 GB	6 GB
• Disk	4.2 GB	21 GB	50.4 GB

GSFC Break-point Analysis (con't.)



Upgrade hardware (16 CPU UltraServers), no new IP addresses:

	<u>Required</u>	<u>Current Design</u>	<u>Hardware Upgrade</u>
Scheduling			
• Job Throughput	32,032	187,200	748,800
• RAM	1.6 GB	2 GB	8 GB
• Disk	4.2 GB	21 GB	50.4+ GB
Planning			
• Time to process a 30 day plan	reasonable	111 min.	28 min.
• RAM	1.6 GB	2 GB	8 GB
• Disk	4.2 GB	21 GB	50.4+ GB

GSFC Break-point Analysis (con't.)



Multiple instances of Planning and Scheduling hardware:

- More than one instance of the UltraServer machines could be added but this would require the use of more than one instance of AutoSys and would require software changes.

LaRC Break-point Analysis



Current hardware fully populated (150 MHz CPUs):

Scheduling	<u>Required</u>	<u>Current Design</u>	<u>Maximum Design</u>
• Job Throughput	9,336	42,120	73,710
• RAM	286 MB	512 MB	512 MB
• Disk	2.8 MB	12.6 GB	50.4 GB
Planning			
• Time to process a 30 day plan	reasonable	96 min.	72 min.
• RAM	286 MB	512 MB	512 MB
• Disk	1.56 GB	12.6 GB	50.4 GB

LaRC Break-point Analysis (con't.)



Upgrade hardware, no new IP addresses:

- **The LaRC hardware would be upgraded to the same configuration as the GSFC hardware upgrade and could then be scaled up (see GSFC Breakpoint Analysis for further information).**

Multiple instances of Planning and Scheduling hardware:

- **More than one instance of the UltraServer machines could be added but this would require the use of more than one instance of AutoSys and would require software changes.**